Fast Benchmarks

Michele Michelotto – INFN Padova
Manfred Alef – GridKa Karlsruhe

Fast Benchmark

- Request mainly from WLCG community via GDB to have a fast benchmark
- Requirements clear
  - Open source
  - Easy to run
  - Small, no download (apart from first download)
- Requirement not clear to me
  - How much fast? Reproducible? Reliable?
  - Single core or multicore?
- Use Cases
  - Run everytime we land on a a queue/VM/Cloud machine?
  - Run to sample the resources available?
  - Run to crosscheck is the HS06 declared are reliable?

An example with Geant4

- Thanks to G.Cosmo and A.Dotti
- Based on Geant4
  - Runs on linux x86-64 and ARM
  - realist description of the geometry of the detector
  - footprint 1/3 to ¼ of real experiment
  - No digitization, no analysis.
  - Cpu bound, no I/O
- Download a bootstrap.sh script from cern
- Running the script download the rest of the program and compile (5 – 10 minutes)
- ./run.sh <numThreads> <numEvents>

LHCB fast benchmark

- New contact with P.Charpentier (LHCb) provided by Manfred
- Small python script running about one minute, single threaded
- Differences in measuring a slot in an idle, half loaded or fully loaded machine

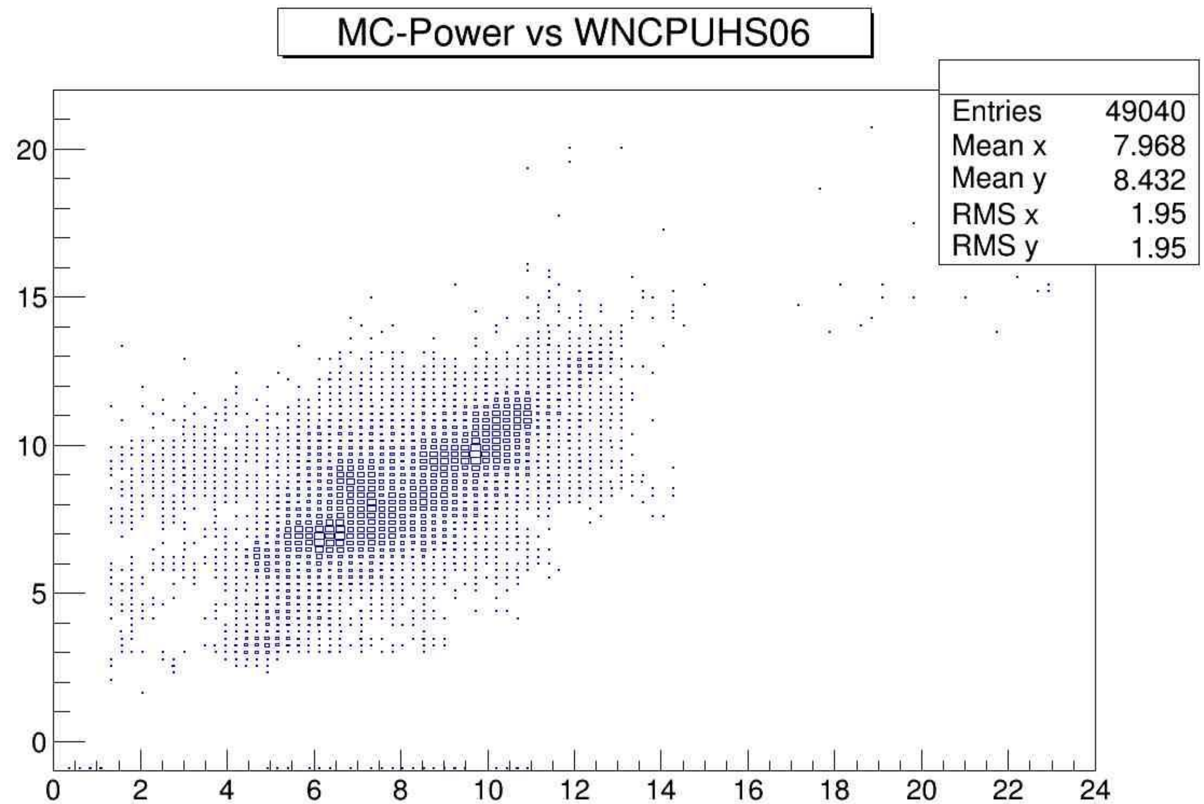

Script from Manfred

- Manfred prepared a script to run several time the LHCB.py script and make averages starting from ONE core loaded to N parallel instances
- N is the number of logical cpus available
- I tried to make a comparison of LHCB.py with HS06 using the load curve of several architectures I had measured in the past

Xeon E5 2660 16C/32T - 2600MHz

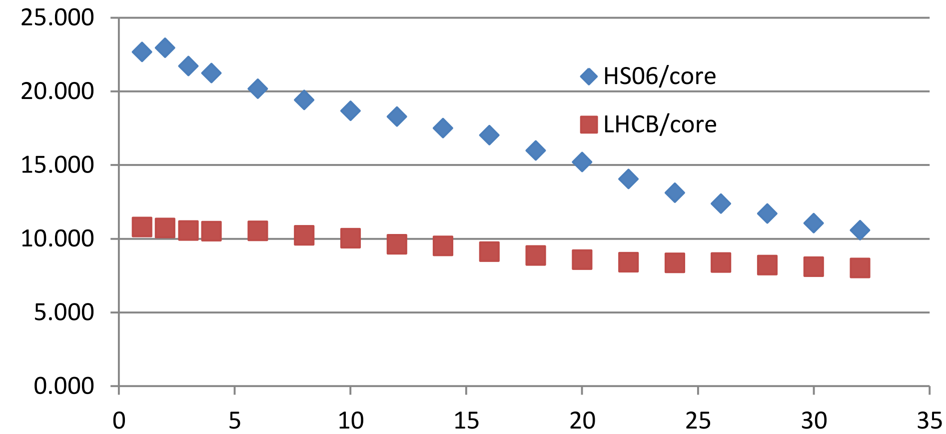

HS06/LHCB

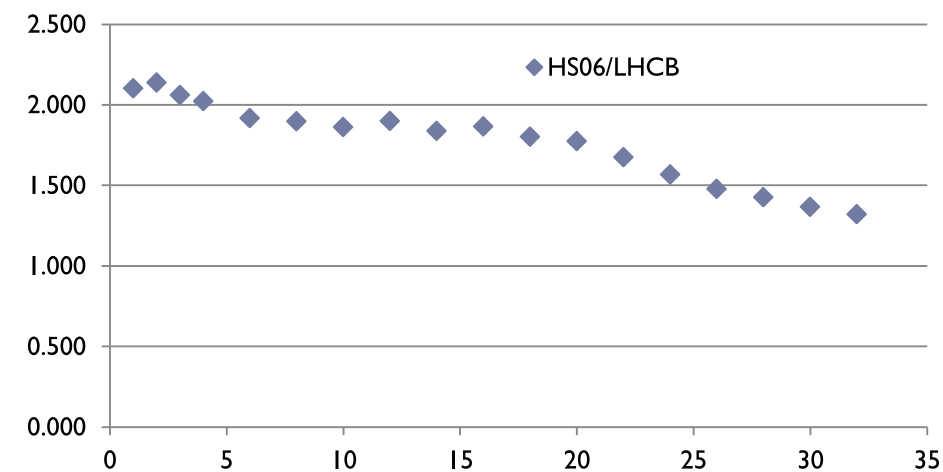

AMD Opteron 6272 (32C - 2100MHz)

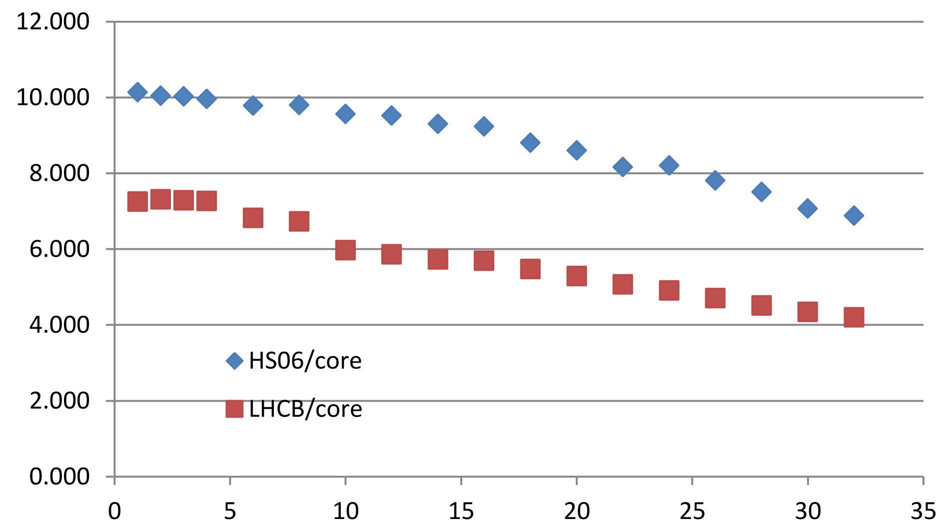

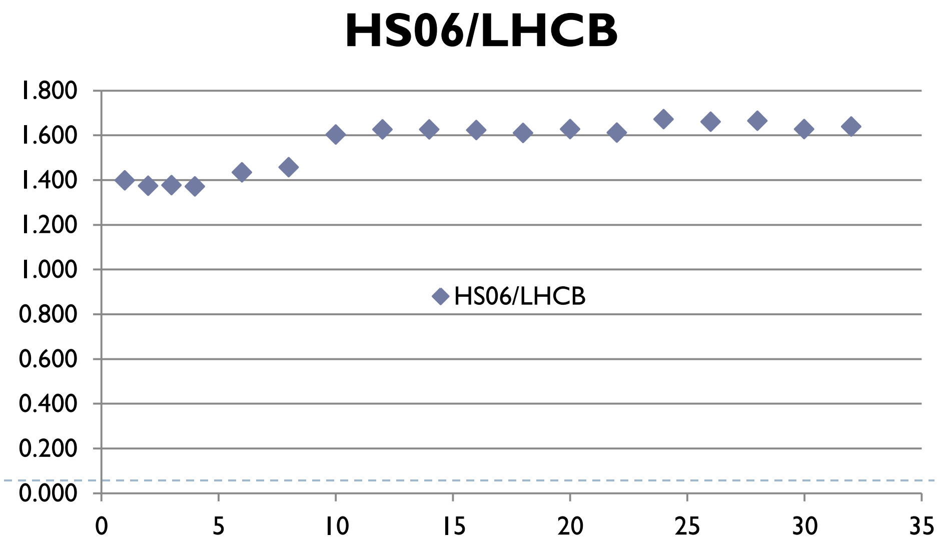

Intel ATOM C2750 (8C @ 2400 MHz)

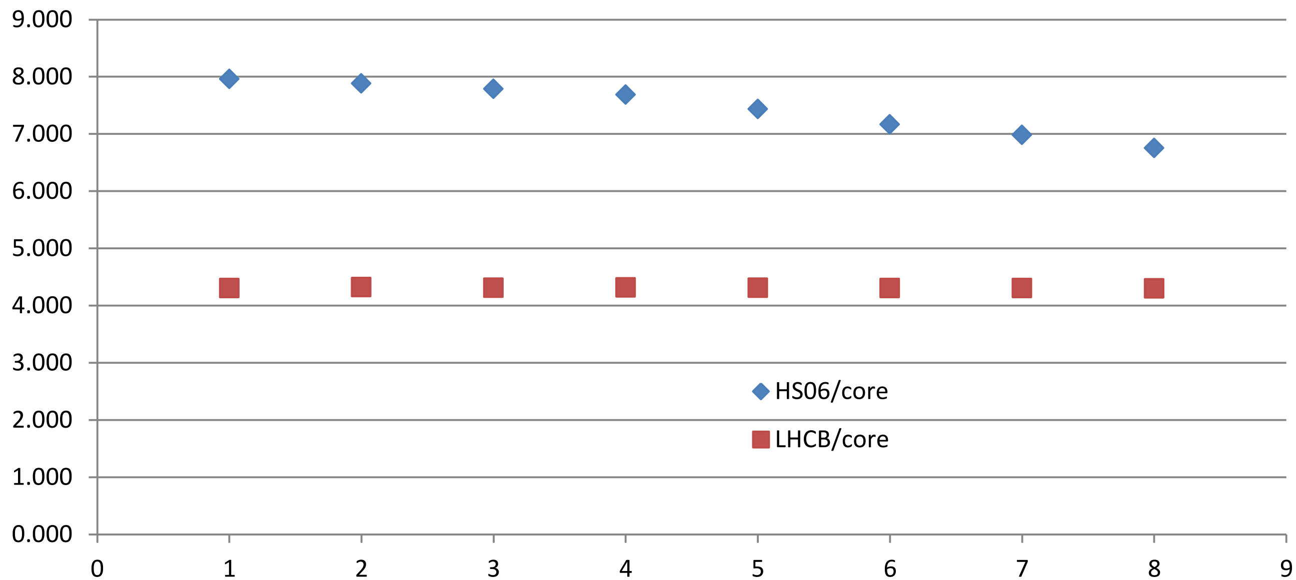

HS06/LHCB

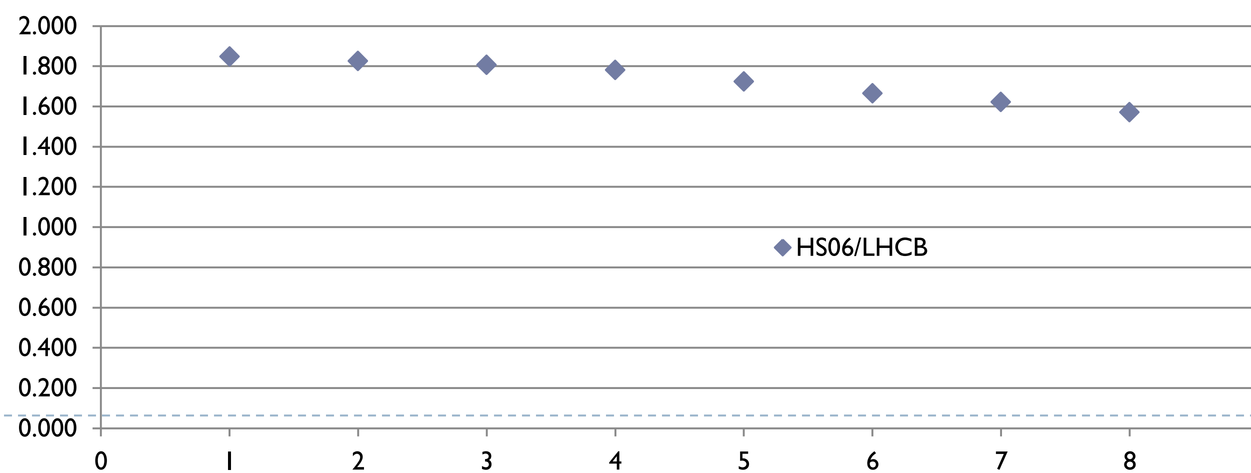

Nvidia Tegra K1 – 4C – 2.3GHz

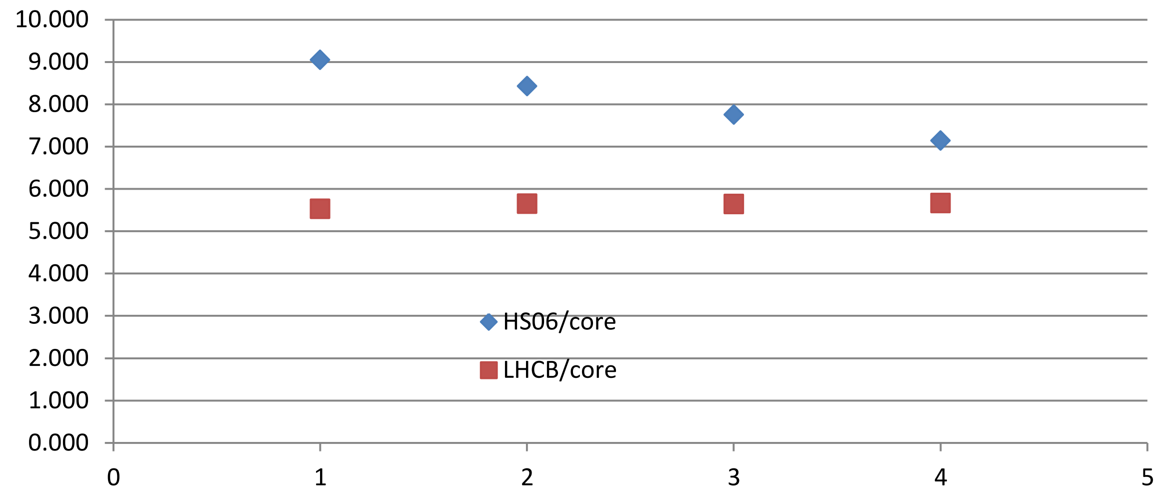

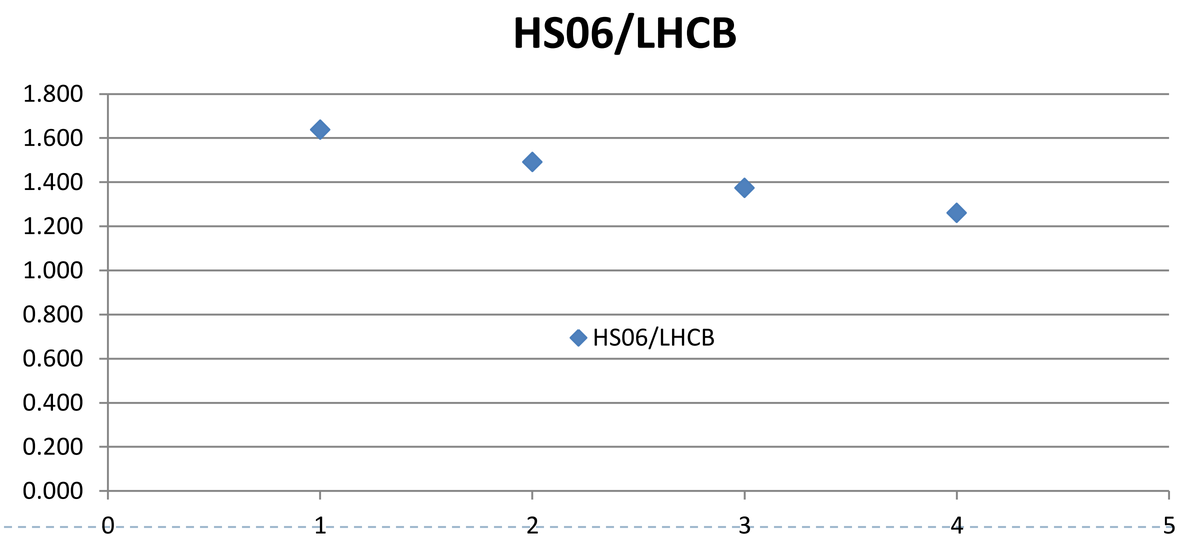

Odroid Exynox 5422 4C - 2.0GHz

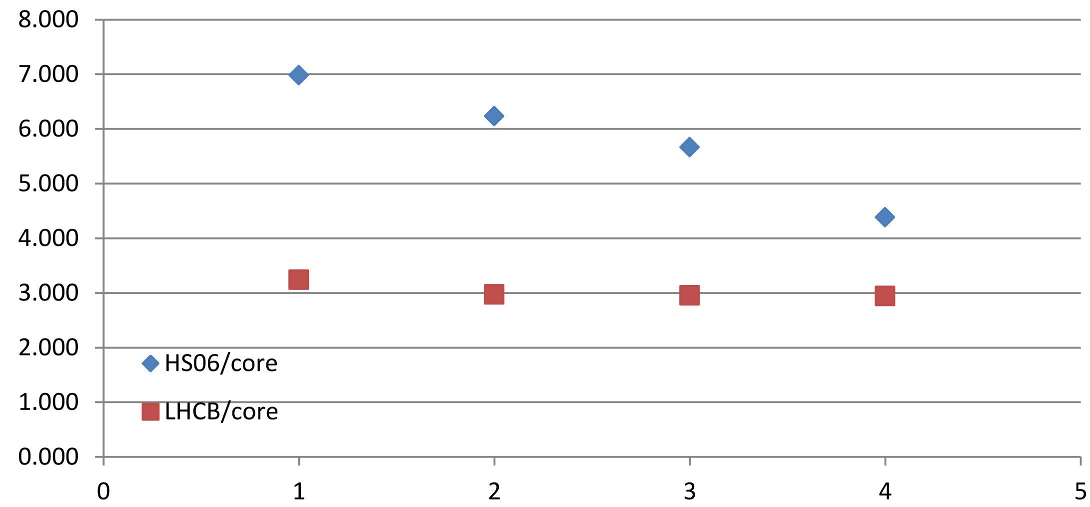

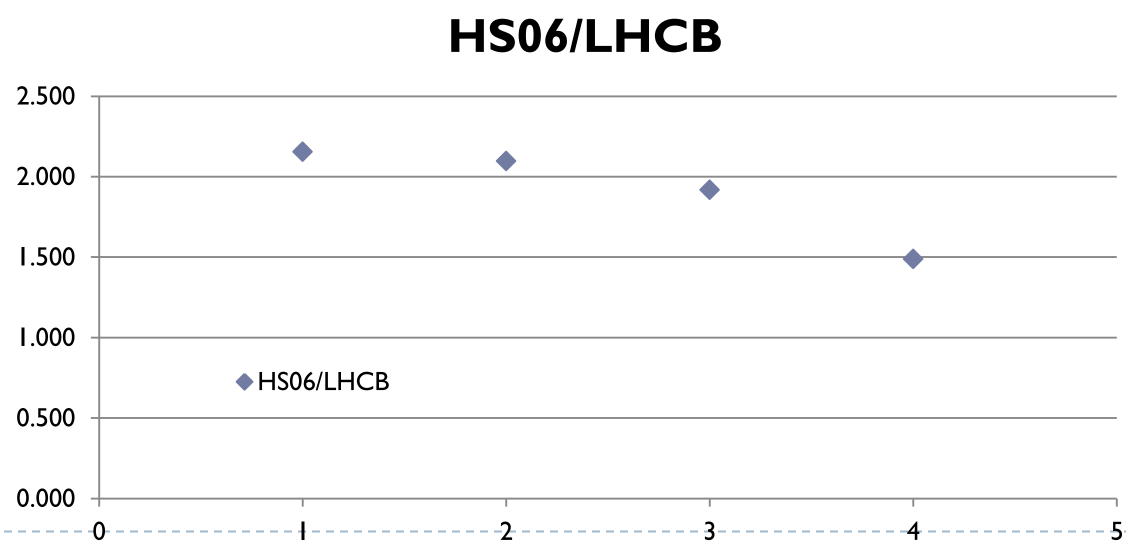

LHCB single instances with HS06 as load

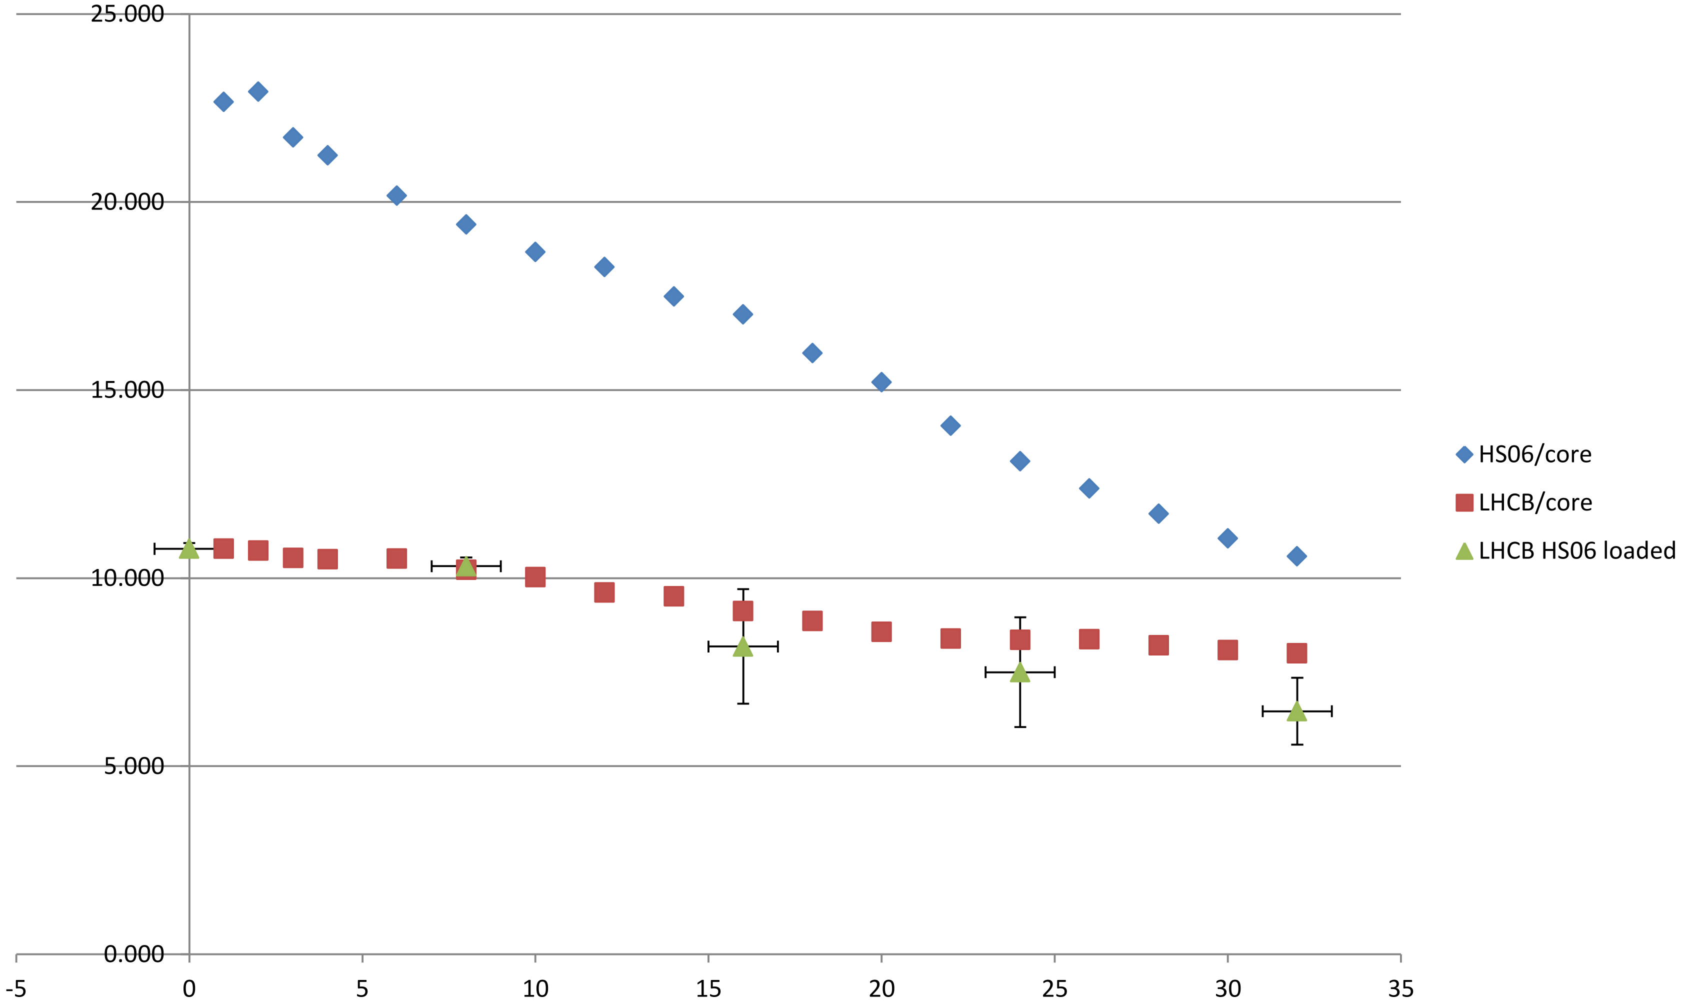

LHCB single instances with HS06 as load

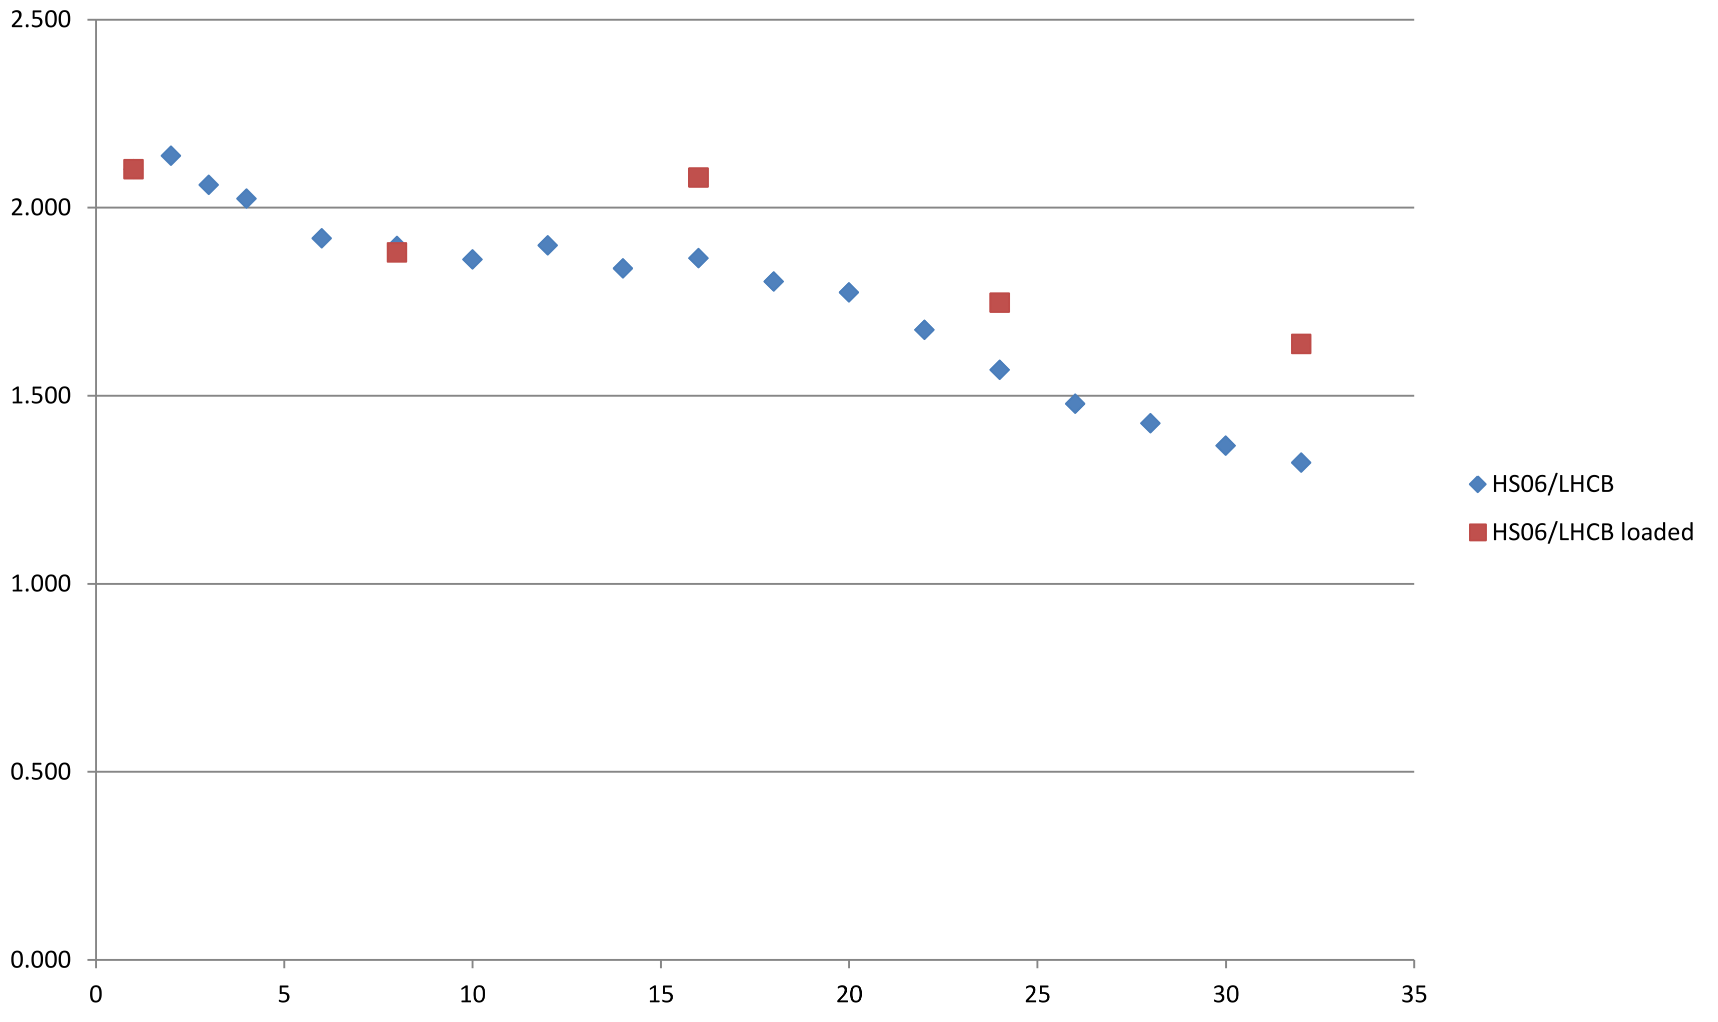

All together now

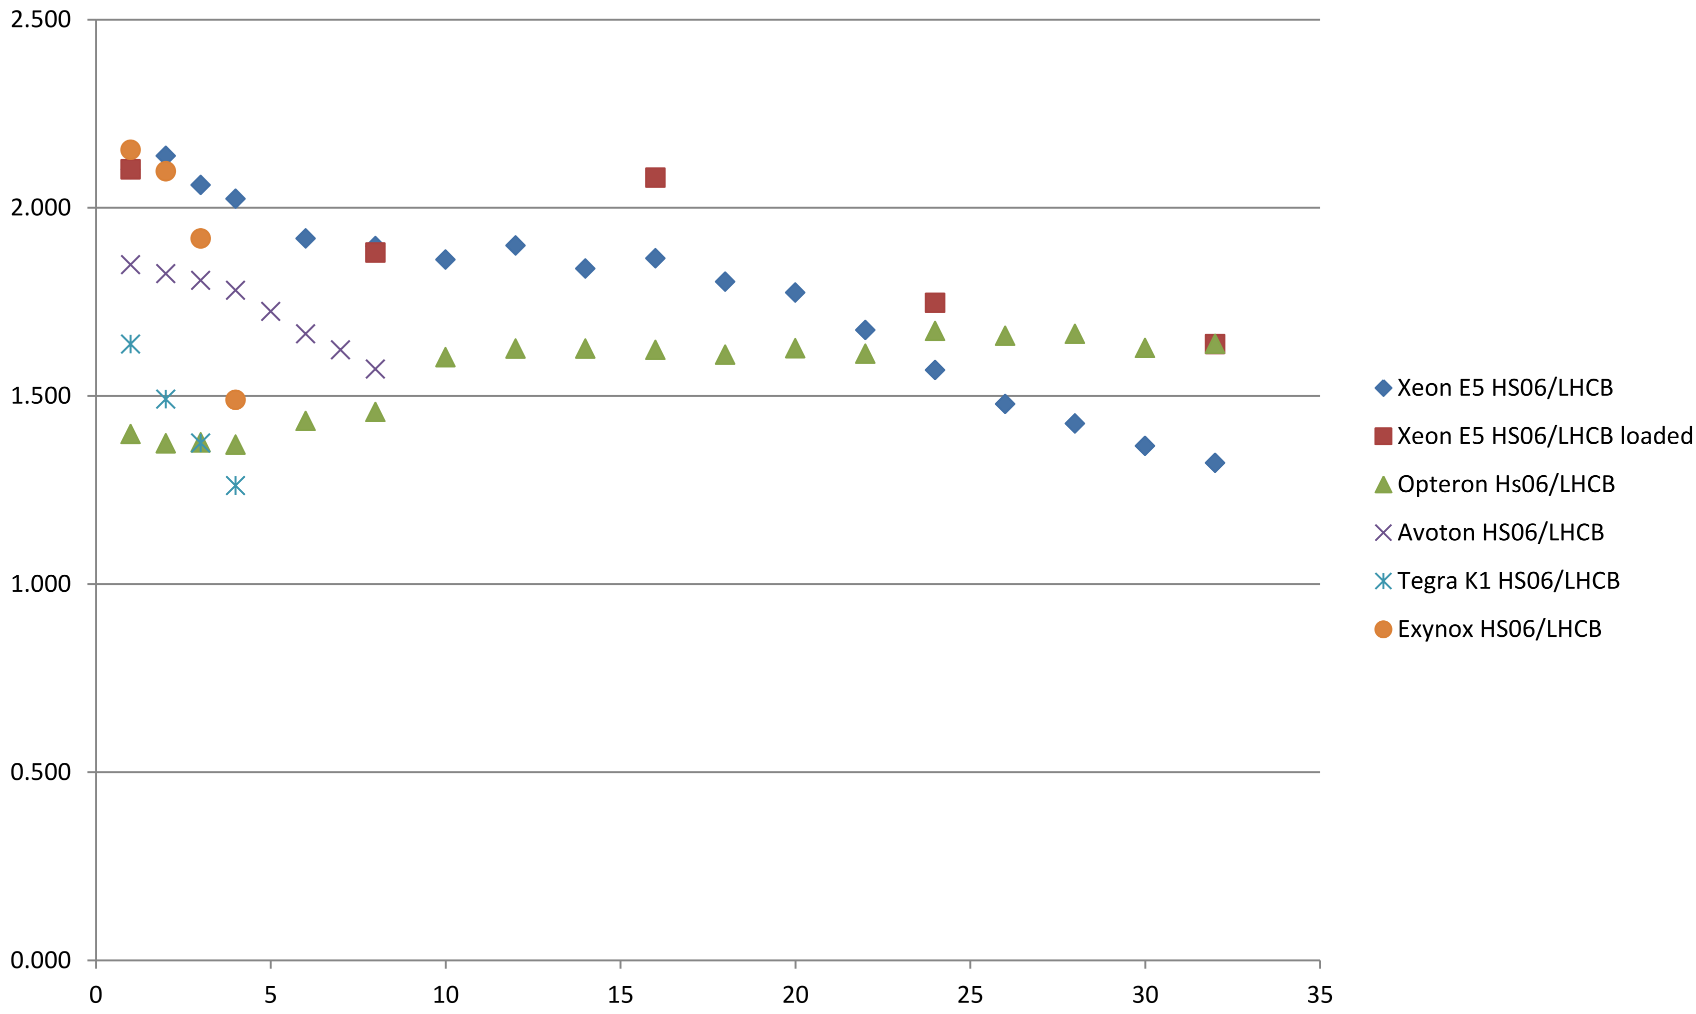

Measuring on a production cluster

- Manfred measured on a GridKa production cluster the LHCb.py score compared with the HS06 per slot
- I made some plot of HS06/LHCB vs load

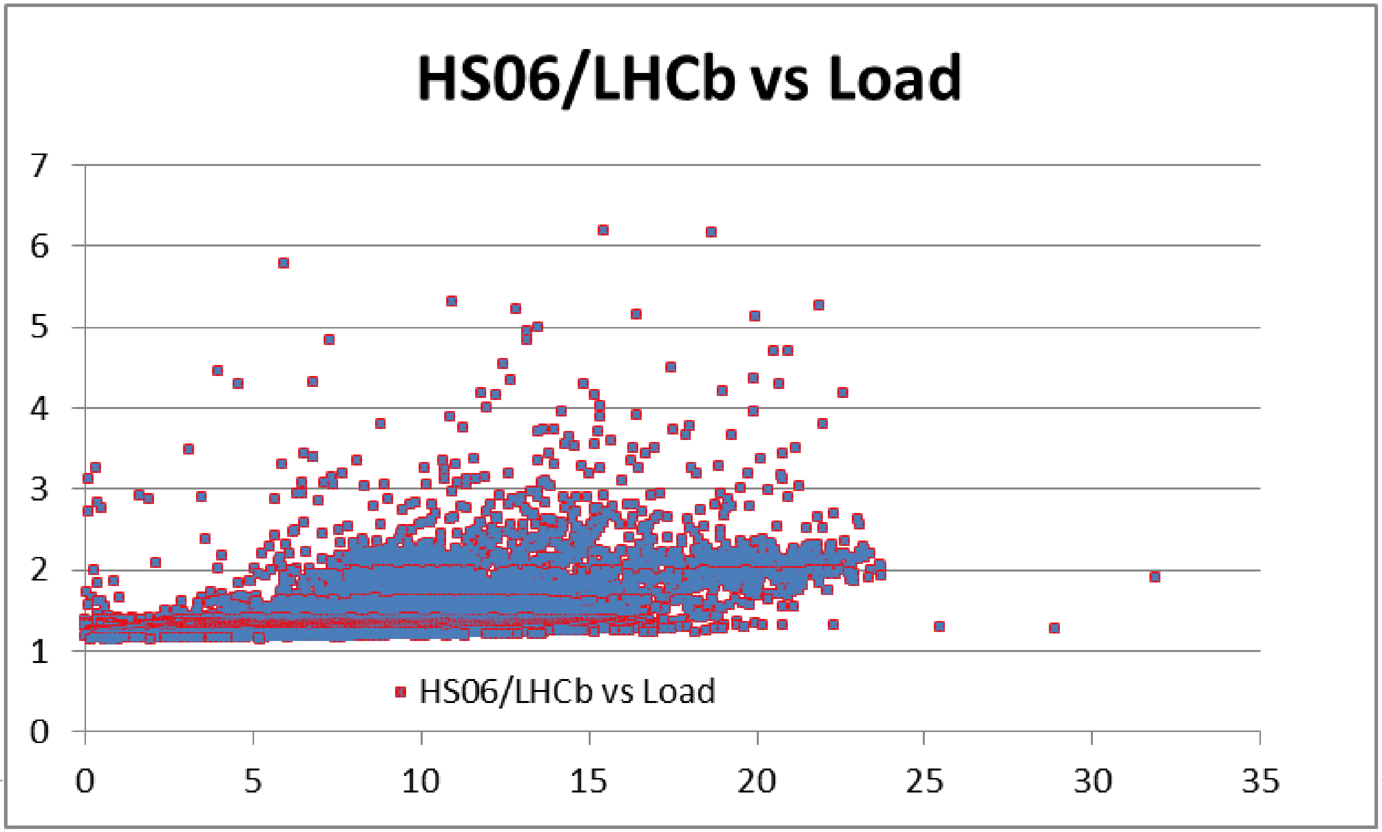

Sometimes LHCb gets slow

- HS06/LHCb score is around 1.2 – 1.6
- Occasionally is can go to more than 2.0

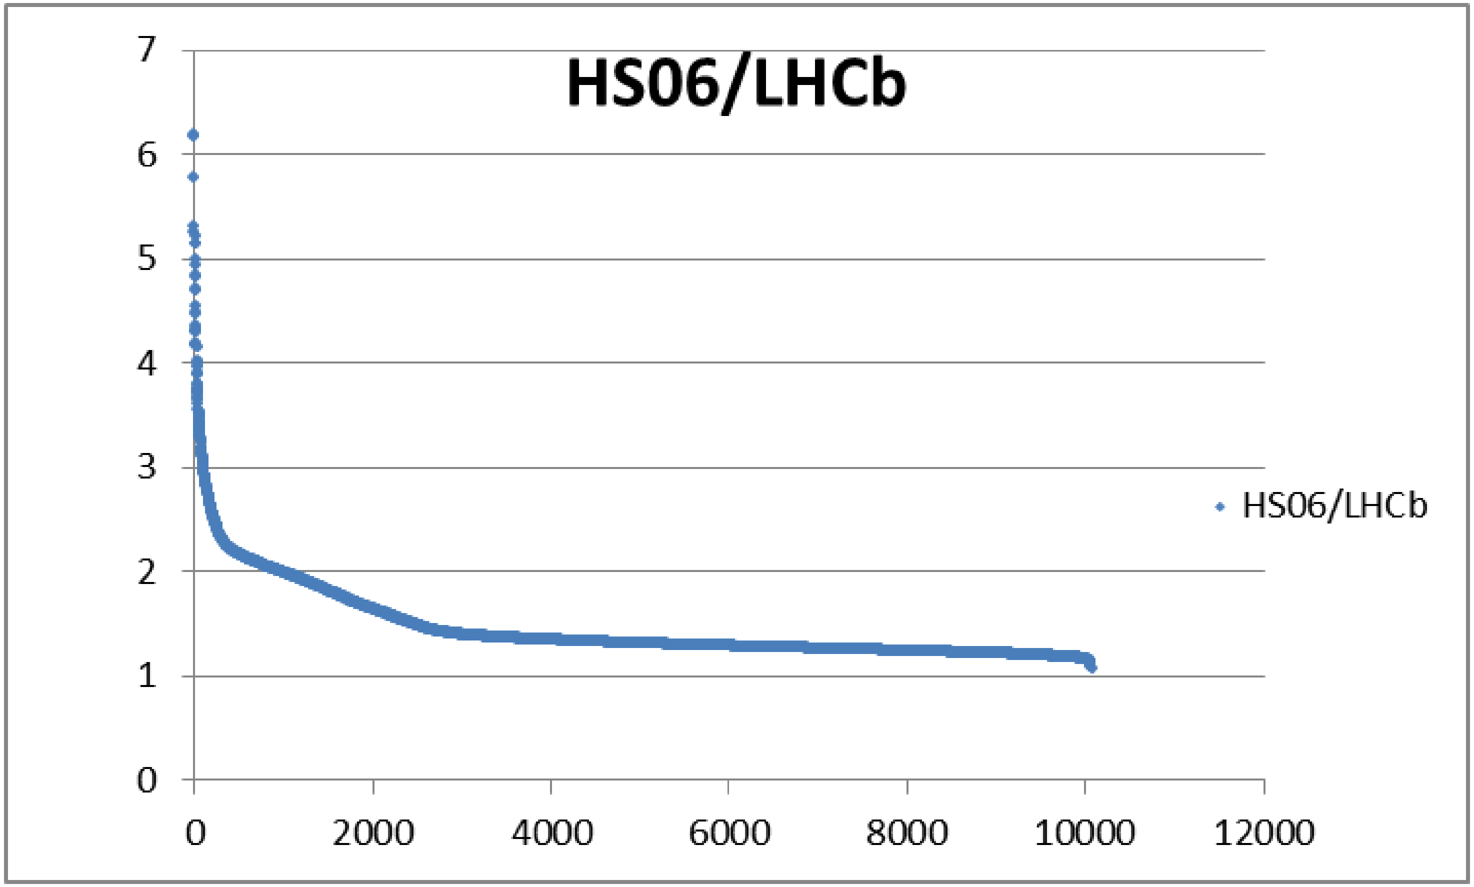

My conclusion

- LHCb.py is a small script that runs easily everywhere you have python
- Is very very simple to maintain. About 30 lines of code
- It takes about one minute to estimate the performance of the cpu on which its running
- The fast answer of course is not free. It's a tradeoff between speed and precision